On the Complexity and Approximability of Budget-Constrained Minimum Cost Flows[☆]

Michael Holzhauser[a,*], Sven O. Krumke[a], Clemens Thielen[a]

[a] *University of Kaiserslautern, Department of Mathematics*
Paul-Ehrlich-Str. 14, D-67663 Kaiserslautern, Germany

arXiv:1607.02282v1 [cs.DS] 8 Jul 2016

Abstract

We investigate the complexity and approximability of the *budget-constrained minimum cost flow problem*, which is an extension of the traditional minimum cost flow problem by a second kind of costs associated with each edge, whose total value in a feasible flow is constrained by a given budget B. This problem can, e.g., be seen as the application of the ε-constraint method to the bicriteria minimum cost flow problem. We show that we can solve the problem exactly in weakly polynomial time $\mathcal{O}(\log M \cdot \mathrm{MCF}(m, n, C, U))$, where C, U, and M are upper bounds on the largest absolute cost, largest capacity, and largest absolute value of any number occuring in the input, respectively, and $\mathrm{MCF}(m, n, C, U)$ denotes the complexity of finding a traditional minimum cost flow. Moreover, we present two fully polynomial-time approximation schemes for the problem on general graphs and one with an improved running-time for the problem on acyclic graphs.

Keywords: algorithms, complexity, minimum cost flow, approximation

1. Introduction

In this paper, we investigate the natural extension of the traditional minimum cost flow problem (cf., e.g., [1]) by a second kind of costs, called *usage fees*, which are linear in the flow on the corresponding edge and bounded by a given budget B. This extension allows us to solve many related problems such as the budget-constrained maximum dynamic flow problem (since each dynamic flow can be represented as a traditional minimum cost flow (cf. [2])) or the application of the ε-constraint method to the bicriteria minimum cost flow problem (cf., e.g., [3]).

To the best of our knowledge, the budget-constrained minimum cost flow problem was first mentioned in [1], where a structural result but no combinatorial algorithm was presented. A related problem in which a *fixed* usage fee is induced by each edge with positive flow was investigated by Duque et al. [4]. The model that we use here was recently investigated in Holzhauser et al. [5], where a strongly polynomial-time algorithm based on the interpretation of the problem as a bicriteria minimum cost flow problem was derived.

We extend these results and show that, using similar ideas, we can also obtain a weakly polynomial-time combinatorial algorithm that performs worse only by a logarithmic factor than the best algorithm for the traditional min-

[☆]This work was partially supported by the German Federal Ministry of Education and Research within the project "SinOptiKom – Cross-sectoral Optimization of Transformation Processes in Municipal Infrastructures in Rural Areas".

[*]Corresponding author. Fax: +49 (631) 205-4737. Phone: +49 (631) 205-2511

Email addresses:
holzhauser@mathematik.uni-kl.de (Michael Holzhauser), krumke@mathematik.uni-kl.de (Sven O. Krumke), thielen@mathematik.uni-kl.de (Clemens Thielen)

imum cost flow problem. Moreover, we present two fully polynomial-time approximation schemes (FPTAS), one of which is based on techniques introduced by Papadimitriou and Yannakakis [6] and has a weakly polynomial running-time and one of which is based on the packing-LP framework developed by Garg and Koenemann [7], achieving a strongly-polynomial running-time. The running-time of the latter FPTAS is subsequently improved for the case of acyclic graphs.

2. Preliminaries

2.1. Problem Definition

In the budget-constrained minimum cost flow problem (abbreviated as $\text{BCMCFP}_{\mathbb{R}}$ in the following), we are given a directed multigraph $G = (V, E)$ with edge capacities $u_e \in \mathbb{N}_{\geq 0}$, costs $c_e \in \mathbb{Z}$ (i.e., we allow integral costs with arbitrary sign), and usage fees $b_e \in \mathbb{N}_{\geq 0}$ per unit of flow on the edges $e \in E$, as well as a budget $B \in \mathbb{N}_{\geq 0}$ and a distinguished source $s \in V$ and sink $t \in V$. The aim is to find a feasible s-t-flow x in G that minimizes $\sum_{e \in E} c_e \cdot x_e$ subject to the budget-constraint $\sum_{e \in E} b_e \cdot x_e \leq B$. The problem $\text{BCMCFP}_{\mathbb{R}}$ can be stated as a linear program as follows:

$$\min \sum_{e \in E} c_e \cdot x_e \tag{1a}$$

$$\text{s.t.} \sum_{e \in \delta^-(v)} x_e - \sum_{e \in \delta^+(v)} x_e = 0 \quad \forall v \in V \setminus \{s, t\}, \tag{1b}$$

$$\sum_{e \in E} b_e \cdot x_e \leq B, \tag{1c}$$

$$0 \leq x_e \leq u_e \qquad \forall e \in E. \tag{1d}$$

Here, we denote by $\delta^+(v)$ ($\delta^-(v)$) the set of *outgoing* (*incoming*) edges of some node $v \in V$. We assume that there are no nodes $v \in V \setminus \{s, t\}$ with $\delta^+(v) = \emptyset$ or $\delta^-(v) = \emptyset$ since no flow can reach such nodes due to flow conservation. Note that we can detect and remove such nodes along with their incident edges in linear time.

Furthermore, note that since the zero-flow is always feasible and has objective value zero, the optimal objective value of each instance of $\text{BCMCFP}_{\mathbb{R}}$ is always non-positive. The problem is a generalization of the problem variant in which a desired flow value F is given since we can "enforce" such a flow value by adding an edge with negative costs of large absolute value and capacity F (cf. [5] for further details).

2.2. Approximation Algorithms

An algorithm A is called a (polynomial-time) *approximation algorithm with performance guarantee* $\alpha \in [1, \infty)$ or simply an *α–approximation* for $\text{BCMCFP}_{\mathbb{R}}$ if, for each instance I of $\text{BCMCFP}_{\mathbb{R}}$ with optimum solution x^*, it computes a feasible solution x with objective value $c(x) \leq \frac{1}{\alpha} c(x^*)$ in polynomial time (note that $c(x)$ and $c(x^*)$ are non-positive, so $\frac{1}{\alpha} c(x^*) \geq c(x^*)$). An algorithm A that receives as input an instance $I \in \Pi$ and a real number $\varepsilon \in (0, 1)$ is called a *polynomial-time approximation scheme (PTAS)* if, on input (I, ε), it computes a feasible solution x with objective value $c(x) \leq (1 - \varepsilon) \cdot c(x^*)$ with a running-time that is polynomial in the encoding size $|I|$ of I. If this running-time is additionally polynomial in $\frac{1}{\varepsilon}$, the algorithm is called a *fully polynomial-time approximation scheme (FPTAS)*.

Moreover, for the case of $\text{BCMCFP}_{\mathbb{R}}$, we call an algorithm a *bicriteria FPTAS* if, for each $\varepsilon \in (0, 1)$, it computes a solution x with $c(x) \leq (1 - \varepsilon) \cdot c(x^*)$ and $b(x) \leq (1 + \varepsilon) \cdot b(x^*)$ in polynomial time.

2.3. Parametric Search

Throughout the paper, we make use of Megiddo's parametric search technique (cf. [8]), which can be described as follows: Assume that we want to solve an optimization problem Π for which we already know an (exact) algorithm A that solves the problem, but in which some of the input values are now *linear parametric values* that depend linearly on some real parameter λ. Moreover, suppose that an algorithm C is known (in the following called *callback*) that is able to decide if some candidate value for λ is smaller, larger, or equal to the value λ^* that leads to an optimum solution to the underlying problem Π. The idea of the parametric search technique is to

simulate the execution of algorithm A with variables that still depend on the symbolic value λ, and to continue the execution until we reach a comparison of two linear parametric values that needs to be resolved. Since both values depend linearly on λ, it either holds that one of the variables is always larger than or equal to the other one (in which case the result of the comparison is independent from λ) or that there is a unique intersection point λ'. For this intersection point, we evaluate the callback C in order to determine if $\lambda' < \lambda^*$, $\lambda' > \lambda^*$, or $\lambda' = \lambda^*$ and, thus, resolve the comparison and continue the execution. Hence, as soon as the simulation of A finishes, we have obtained an optimum solution to Π. The overall running-time is given by the running-time of A times the running-time of C and can be further improved using parallelization techniques described in [9]. We refer to [8] for further details on the parametric search technique. Further applications and extenions of parametric search techniques can moreover be found in [10, 11, 12].

3. Exact Algorithms

We start with results on the complexity of the problem BCMCFP$_{\mathbb{R}}$. The mathematical model (1a) – (1c) for BCMCFP$_{\mathbb{R}}$, as introduced in Section 2, is a linear program, which can be solved in weakly polynomial time by known techniques such as interior point methods (cf. [13]). In particular, using the procedure described by Vaidya [14] to our multigraph setting, we get the following weakly polynomial running-time for BCMCFP$_{\mathbb{R}}$:

Theorem 1. BCMCFP$_{\mathbb{R}}$ *is solvable in weakly polynomial time* $\mathcal{O}(m^{2.5} \cdot \log M)$. □

However, in this paper, we are interested in *combinatorial* algorithms that exploit the structure of the underlying problem. We show how we can incorporate combinatorial algorithms for the traditional minimum cost flow problem in order to solve the more general budget-constrained minimum cost flow problem BCMCFP$_{\mathbb{R}}$.

We can solve BCMCFP$_{\mathbb{R}}$ by computing an efficient solution of the bicriteria minimum cost flow problem with the two objective functions $c(x)$ and $b(x)$ that minimizes $c(x)$ while maintaining $b(x) \leq B$. Graphically, each optimum solution x^* of BCMCFP$_{\mathbb{R}}$ corresponds to a point in the objective space that lies on the pareto frontier and not above the line $b = B$.[1] The situation is shown in Figure 1.

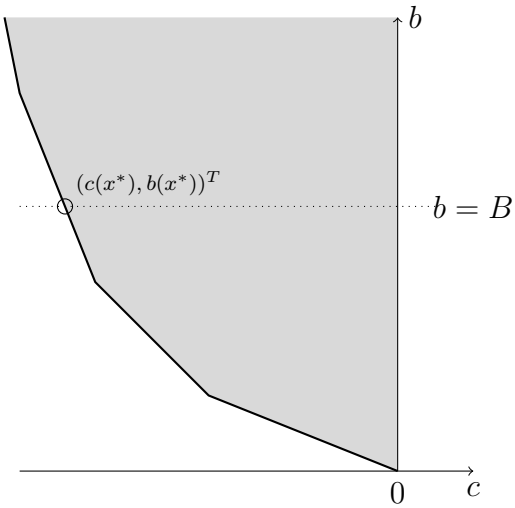

Figure 1: The objective space of the interpretation of BCMCFP$_{\mathbb{R}}$ as a bicriteria minimum cost flow problem. The gray area corresponds to the set of the objective values of feasible flows, the thick black lines correspond to efficient edges, which form the pareto frontier.

It is well-known that, for each point $(c, b)^T$ on the pareto frontier, there is some value $\lambda \in [0, \infty)$ and a feasible flow x with $(c(x), b(x))^T = (c, b)^T$ such that x is a minimum cost flow with respect to the costs $b_e + \lambda \cdot c_e$ for each edge $e \in E$ (cf. Geoffrion [16]). Assume that there are two flows x^1 and x^2 that are both optimal for some specific value of λ, i.e., $b(x^1)+\lambda \cdot c(x^1) = b(x^2)+\lambda \cdot c(x^2) = \alpha$ for some value α. Then, for both of the flows x^i with $i \in \{1, 2\}$, it holds that $b(x^i) = \alpha - \lambda \cdot c(x^i)$, i.e., they lie on the same efficient edge, which is a straight line with slope $-\lambda$ in the objective space. In other words, computing a minimum cost flow with edge-costs $b_e + \lambda \cdot c_e$ will either provide a solution that corresponds to on an extreme point of the pareto frontier or some point that lies on the efficient edge with slope $-\lambda$. Moreover, as shown

[1] We refer to [15] for an in-depth treatment of bicriteria optimization problems and efficient solutions.

in [5], the slopes of these efficient edges differ by minimum absolute amounts:

Lemma 1 ([5]). *The slopes of two efficient edges on the pareto frontier of any instance of* $\text{BCMCFP}_{\mathbb{R}}$ *differ by an absolute value of at least* $\frac{1}{\overline{c}^2}$ *for* $\overline{c} := \sum_{e \in E} |u_e \cdot c_e|$. □

As explained above, each optimum solution x^* of $\text{BCMCFP}_{\mathbb{R}}$ is a minimum cost flow with respect to the edge-costs $c_e + \lambda^* \cdot b_e$ for at least one value $\lambda^* \in [0, +\infty)$. In particular, if Λ^* denotes the set of all these values λ^*, it holds that Λ^* is a closed interval containing either one or infinitely many such values λ^* depending on whether the optimum solutions correspond to points that lie amid or at the corner of some efficient edge in the objective space, respectively. As claimed in the following lemma, which is proven in [5], we are able to decide the membership in Λ^* efficiently:

Lemma 2 ([5]). *Let* $\Lambda^* \neq \emptyset$ *denote the set of parameters* λ^* *for which an optimum solution* x^* *to* $\text{BCMCFP}_{\mathbb{R}}$ *is a minimum cost flow with respect to the edge-costs* $c_e + \lambda^* \cdot b_e$ *for each* $e \in E$. *For some candidate value* λ, *it is possible to decide whether* $\lambda < \min \Lambda^*$, $\lambda > \max \Lambda^*$, *or* $\lambda \in \Lambda^*$ *in* $\mathcal{O}(\text{MCF}(m, n, C, U))$ *time.* □

Let $\text{MCF}(m, n)$ denote the complexity of computing a traditional minimum-cost flow in *strongly* polynomial time. As shown in [5], Lemma 2 can be used within Megiddo's parametric search technique in order to obtain strongly polynomial-time algorithms for $\text{BCMCFP}_{\mathbb{R}}$:

Theorem 2 ([5]). $\text{BCMCFP}_{\mathbb{R}}$ *is solvable in strongly polynomial time* $\mathcal{O}(m \log m \cdot \min\{T_1(m, n), T_2(m, n), T_3(m, n)\})$ *with*

- $T_1(m, n) \in \mathcal{O}((m + n \log n) \cdot \text{MCF}(m, n))$,
- $T_2(m, n) \in \mathcal{O}((n \cdot \log \frac{m}{n} + n \cdot \log n + \log m) \cdot \text{MCF}(m, n) + m)$, *and*
- $T_3(m, n) \in \mathcal{O}(\log \log m \cdot \log^2 m \cdot \text{MCF}(m, n) + f(m))$ *for* $f(m) \in o(m^3)$. □

We now show how we can use Lemma 2 within a binary search in order to obtain a weakly polynomial-time algorithm that performs within a factor $\mathcal{O}(\log M)$ of each algorithm for the traditional minimum cost flow problem:

Theorem 3. $\text{BCMCFP}_{\mathbb{R}}$ *is solvable in weakly polynomial time* $\mathcal{O}(\log M \cdot \text{MCF}(m, n, C, U))$.

Proof. Consider the set $K := \left\{k \cdot \frac{1}{2\overline{c}^2} : k \in \{0, \ldots, \overline{b} \cdot 2\overline{c}^2\}\right\}$, where $\overline{b} := \sum_{e \in E} u_e \cdot b_e \in \mathcal{O}(m \cdot M^2)$ and $\overline{c} := \sum_{e \in E} |u_e \cdot c_e| \in \mathcal{O}(m \cdot M^2)$ are upper bounds on the total usage fees and total absolute value of the costs of any feasible flow, respectively. Note that each extreme point of the pareto frontier can be obtained by a minimum cost flow computation with edge-costs $c_e + \lambda \cdot b_e$ for some $\lambda \in K$ since the slopes of any two efficient edges differ by an absolute amount of at least $\frac{1}{\overline{c}^2}$ according to Lemma 1. Hence, by incorporating the procedure that is described in Lemma 2 into a binary search on the set K, we either find some value $\lambda \in \Lambda^*$ (in which case we have also found an optimum solution to $\text{BCMCFP}_{\mathbb{R}}$) or two "adjacent" values $\lambda^{(1)} := k \cdot \frac{1}{2\overline{c}^2}$ and $\lambda^{(2)} := (k+1) \cdot \frac{1}{2\overline{c}^2}$ for some $k \in \{0, \ldots, \overline{b} \cdot 2\overline{c}^2 - 1\}$ with $\lambda^{(1)} < \min \Lambda^*$ and $\lambda^{(2)} > \max \Lambda^*$. These values, however, yield solutions $x^{(1)}$ and $x^{(2)}$ that correspond to the corner points of the same efficient edge, which crosses the line $b = B$ in the objective space. Thus, by computing a suitable convex combination of the two solutions $x^{(1)}$ and $x^{(2)}$, we obtain an optimum solution to $\text{BCMCFP}_{\mathbb{R}}$ and are done.

The running-time of the procedure is dominated by the binary search on the set K and the resulting $\mathcal{O}(\log |K|)$ calls to the procedure that is described in Lemma 2. Hence, the overall running-time is given by

$$\begin{aligned} &\mathcal{O}(\log |K| \cdot \text{MCF}(m, n, C, U)) \\ &= \mathcal{O}(\log(\overline{b} \cdot \overline{c}^2) \cdot \text{MCF}(m, n, C, U)) \\ &= \mathcal{O}(\log(m^3 \cdot M^6) \cdot \text{MCF}(m, n, C, U)) \\ &= \mathcal{O}(\log M \cdot \text{MCF}(m, n, C, U)), \end{aligned}$$

which shows the claim. □

4. Approximation Algorithms

4.1. General Graphs

In [6], the authors show that an *ε-approximate pareto frontier* (i.e., a set of points P_ε such that, for each point y on the pareto frontier P, there is a point $y' \in P_\varepsilon$ such that y is within a factor of $(1+\varepsilon)$ from y' in each component) of a linear convex optimization problem with k objective functions can be determined by solving $\mathcal{O}((8Lk^2/\varepsilon)^k)$ instances of the problem with only one objective function, where L denotes the encoding-length of the largest possible objective value. Applied to $\text{BCMCFP}_{\mathbb{R}}$, we are, thus, able to determine an ε-approximate pareto frontier in $\mathcal{O}\left(\left(\frac{\log M}{\varepsilon}\right)^2 \cdot \text{MCF}(m,n,C,U)\right)$ time, which in turn implies a bicriteria FPTAS for $\text{BCMCFP}_{\mathbb{R}}$. The following lemma shows that this also yields a traditional FPTAS for $\text{BCMCFP}_{\mathbb{R}}$:

Lemma 3. *Any bicriteria FPTAS for* $\text{BCMCFP}_{\mathbb{R}}$ *also induces a single-criterion FPTAS for* $\text{BCMCFP}_{\mathbb{R}}$.

Proof. For each instance of $\text{BCMCFP}_{\mathbb{R}}$ with optimum solution x^*, the given bicriteria FPTAS computes a solution x with $c(x) \leq (1-\varepsilon)\cdot c(x^*)$ and $b(x) \leq (1+\varepsilon)\cdot b(x^*)$ in time that is polynomial in the instance size and $\frac{1}{\varepsilon}$. Since both the costs c and usage fees b are linear functions, it suffices to scale down the given solution as follows: Let $x' := \frac{x}{1+\varepsilon}$. Clearly, x' is feasible since it still fulfills every flow conservation and capacity constraint and since $b(x') = \frac{1}{1+\varepsilon}\cdot b(x) \leq b(x^*) \leq B$. Moreover, for $\varepsilon' := 2\varepsilon$, it holds that

$$\begin{aligned} c(x') &= \frac{1}{1+\varepsilon}\cdot c(x) \leq \frac{1-\varepsilon}{1+\varepsilon}\cdot c(x^*) \\ &= \frac{1-\frac{\varepsilon'}{2}}{1+\frac{\varepsilon'}{2}}\cdot c(x^*) = \frac{1-\varepsilon'+\frac{(\varepsilon')^2}{4}}{1-\frac{(\varepsilon')^2}{4}}\cdot c(x^*) \\ &\leq (1-\varepsilon')\cdot c(x^*), \end{aligned}$$

which shows the claim. □

Corollary 1. *There is an FPTAS for* $\text{BCMCFP}_{\mathbb{R}}$ *that runs in* $\mathcal{O}\left(\left(\frac{\log M}{\varepsilon}\right)^2 \cdot \text{MCF}(m,n,C,U)\right)$ *time.* □

We now show how we can obtain an FPTAS with *strongly* polynomial running-time using a different approach based on a combination of Garg and Koenemann's packing-LP framework and Megiddo's parametric search technique. We will therefore need the following auxiliary lemma:

Lemma 4. *Let λ be a parameter with a callback that fulfills $C(m,n) \in \Omega(\frac{m}{\log m})$. Any multigraph G with linear parametric edge-lengths $l_e(\lambda)$ on each $e \in E$ can be turned into a simple graph G' that only contains the shortest edge among all parallel edges between two nodes in $\mathcal{O}(\log m \cdot \log\log m \cdot C(m,n))$ time.*

Proof. Let $S := \{(v,w) \in V^2 : |\delta^+(v) \cap \delta^-(w)| \geq 2\}$ denote the set of all pairs of nodes with at least two parallel edges between them. In order to determine the simple graph G' with the desired properties, we need to evaluate the minimum of all edges in $\delta^+(v) \cap \delta^-(w)$ for each $(v,w) \in S$. As shown in [17], we can determine the minimum of k values in $\mathcal{O}(\log\log k)$ time using $\mathcal{O}(k)$ processors. We simulate all of these computations in parallel, which results in a total number of $\mathcal{O}(\sum_{(v,w)\in S} |\delta^+(v) \cap \delta^-(w)|) = \mathcal{O}(m)$ processors. In order to reduce the number of callback calls, we simulate the $\mathcal{O}(m)$ processors sequentially in a round-robin manner until each of them either finishes its computation or holds at the comparison of two linear parametric values, yielding $\mathcal{O}(m)$ candidate values for λ that need to be resolved using the callback for λ. Using a binary search on the set of these candidate values in combination with a successive determination of the median, which can be employed in $\mathcal{O}(m)$ time according to Blum et al. [18], we can resolve all of the comparisons simultaneously in $\mathcal{O}(\log m \cdot C(m,n) + m)$ time and continue the simulation of the processors. After $\mathcal{O}(\log\log m)$ iterations of the above procedure, each processor has finished its computation and the edge with minimum length is determined for each $(v,w) \in S$, which shows the claim. □

Theorem 4. *There is an FPTAS for* $\text{BCMCFP}_{\mathbb{R}}$ *that runs in strongly polynomial-time* $\widetilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2}\cdot(m^2\cdot n + m\cdot n^3)\right)$.[2]

PROOF. We consider an equivalent, circulation-based version of $\text{BCMCFP}_{\mathbb{R}}$ that can be obtained by inserting an edge with infinite capacity, zero costs, and zero usage fees between t and s. Then, according to the flow decomposition theorem for traditional flows (cf. [1]), each optimum flow x^* is positive on $\mathcal{O}(m)$ simple cycles C with strictly negative costs $c(C) := \sum_{e \in C} c_e < 0$. Let $\mathcal{C}$ denote the set of all such simple cycles with negative costs the underlying graph. For $b(C) := \sum_{e \in C} b_e$, we obtain the following cycle-based formulation of $\text{BCMCFP}_{\mathbb{R}}$:

$$\begin{aligned}
\min \ & \sum_{C \in \mathcal{C}} c(C) \cdot x_C & & \\
\text{s.t.} \ & \sum_{\substack{C \in \mathcal{C}: \\ e \in C}} x_C \leq u_e & & \forall e \in E, \\
& \sum_{C \in \mathcal{C}} b(C) \cdot x_C \leq B, & & \\
& x_C \geq 0 & & \forall C \in \mathcal{C}.
\end{aligned}$$

The dual of this linear program can be stated as follows:

$$\begin{aligned}
\min \ & B \cdot \mu + \sum_{e \in E} u_e \cdot y_e, & & \\
\text{s.t.} \ & b(C) \cdot \mu + \sum_{e \in C} y_e \geq -c(C) & & \forall C \in \mathcal{C}, \quad (2) \\
& y_e \geq 0 & & \forall e \in E, \\
& \mu \geq 0. & &
\end{aligned}$$

Although the number of variables (constraints) is exponential in the primal (dual), we are able to derive a strongly polynomial-time FPTAS for the problem using the packing-LP framework introduced by Garg and Koenemann [7]. In short, Garg and Koenemann show that, as long as it is possible to determine the *most-violated constraint* of the dual for some infeasible dual solution (y, μ) with $y > 0$ and $\mu > 0$ in polynomial time $\mathcal{O}(A(m,n))$, then there is an FPTAS with a running-time in $\mathcal{O}(\frac{1}{\varepsilon^2} \cdot m \log m \cdot A(m,n))$.

[2]To simplify running-times, it is common to use $\widetilde{\mathcal{O}}(p)$ in order to denote $\mathcal{O}(p \cdot \log^k m)$ with $k \in \mathcal{O}(1)$.

Note that, since $c(C) < 0$ for each $C \in \mathcal{C}$, we can rewrite equation (2) as

$$\frac{b(C) \cdot \mu + \sum_{e \in C} y_e}{-c(C)} \geq 1,$$

or, equivalently,

$$\frac{\sum_{e \in C} (b_e \cdot \mu + y_e)}{\sum_{e \in C} -c_e} \geq 1.$$

Hence, we are done if we can determine the *minimum ratio cycle* $C \in \mathcal{C}$ with edge-costs $b_e \cdot \mu + y_e$ and edge-times $-c_e$ in polynomial time. Megiddo [9] derived an algorithm that determines the minimum ratio cycle in a *simple* graph in $\mathcal{O}(n^3 \log n + m \cdot n \log^2 n \log \log n)$ time by computing all-pair shortest paths in combination with Karp's minimum mean cycle algorithm [19] as a callback function in his parametric search. In order to comply with our setting of multigraphs, we first need to apply Lemma 4 with the minimum mean cycle algorithm as a callback, running in $C(m,n) := \mathcal{O}(m \cdot n)$ time, to the underlying graph, which yields a running-time of $\mathcal{O}(m \cdot n \cdot \log m \cdot \log \log m)$. In total, we get that

$$\begin{aligned}
A(m,n) = \ & \mathcal{O}(m \cdot n \cdot \log m \log \log m \\
& + n^3 \log n + m \cdot n \cdot \log^2 n \log \log n).
\end{aligned}$$

Thus, incorporated in Garg and Koenemann's framework, we obtain an overall running-time of

$$\begin{aligned}
& \mathcal{O}\left(\frac{1}{\varepsilon^2} \cdot m \log m \cdot A(m,n)\right) \\
= \ & \widetilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2} \cdot (m^2 \cdot n + m \cdot n^3)\right).
\end{aligned}$$

□

4.2. Acyclic Graphs

We now show how we can improve the running-time of the FPTAS described in Theorem 4 for the case of an acyclic graph G. Since there are no cycles in G, we only need to repeatedly determine minimum ratio s-t-paths rather than minimum ratio cycles. This, however, can be done more efficiently, as shown in the following lemma:

Lemma 5. *Let $d^{(1)}\colon E \to \mathbb{R}$ and $d^{(2)}\colon E \to \mathbb{R}$ be two cost functions with $d_e^{(1)} := d^{(1)}(e)$ and $d_e^{(2)} := d^{(2)}(e)$ for each $e \in E$ and assume that $\sum_{e\in P} d_e^{(2)} > 0$ for each s-t-path P. An s-t-path P^* that minimizes the ratio $\frac{\sum_{e\in P} d_e^{(1)}}{\sum_{e\in P} d_e^{(2)}}$ among all s-t-paths P can be found in $\mathcal{O}(m \cdot (\log m \log\log m + n \cdot \log n))$ time on acyclic graphs.*

Proof. Let $\mathcal{P}$ denote the set of all s-t-paths in the underlying acyclic graph G. Similar to Megiddo [8] and Chandrasekaran [20], we can restrict our considerations to the problem $\min_{P\in\mathcal{P}} \sum_{e\in P} d_e^{(\lambda)}$ with $d_e^{(\lambda)} := d_e^{(1)} - \lambda \cdot d_e^{(2)}$ for each $e \in E$ and some parameter λ: Using similar arguments as in [20], it is easy to see that, for some candidate value of λ, it holds that $\min_{P\in\mathcal{P}} \sum_{e\in P} d_e^{(\lambda)}$ is negative (positive) if and only if the value of λ is smaller (larger) than the value λ^* that leads to an optimum solution P^* to $\min_{P\in\mathcal{P}} \frac{\sum_{e\in P} d_e^{(1)}}{\sum_{e\in P} d_e^{(2)}}$. Hence, by simulating the shortest path algorithm for acyclic graphs with edge lengths $d^{(\lambda)}$ for a symbolic value of λ using Megiddo's parametric search technique, it is possible to determine the optimum solution P^* in $\mathcal{O}(m^2)$ time.

We can improve this running-time by first applying Lemma 4 to the underlying multi-graph in order to obtain a simple graph with the same shortest paths as in G in $\mathcal{O}(m \cdot \log m \log\log m)$ time. In this simple graph, we simulate the shortest path algorithm for acyclic graphs, which initially sets the distance label of each node to infinity. It then investigates the nodes in the order of a topological sorting and, for each outgoing edge $e = (v, w)$ of some node $v \in V \setminus \{t\}$ in this sorting, updates the distance label $dist(w)$ of node w to $\min\{dist(w), dist(v)+l_e\}$ where l_e denotes the length of edge e, which results in a comparison of two linear parametric values. Note that the edges in $\delta^+(v)$ head to different nodes since the underlying graph is simple, so all of these comparisons are independent from each other. Thus, by evaluating a binary search over the set of candidate values that result from each of these comparisons as described in [9] and as used above, we only need $\mathcal{O}(\log|\delta^+(v)|)$ shortest path computations at an overhead of $\mathcal{O}(|\delta^+(v)|)$ for the median computations. This results in an running-time for the parametric shortest path computation of

$$\begin{aligned}&\mathcal{O}\left(\sum_{v\in V\setminus\{t\}} (\log|\delta^+(v)| \cdot m + |\delta^+(v)|)\right)\\ &= \mathcal{O}\left(m \cdot n \cdot \log n\right),\end{aligned}$$

which in, combination with the overhead of $\mathcal{O}(m \cdot \log m \cdot \log\log m)$ for the transformation into a simple graph, shows the claim. □

By incorporating the results of Lemma 5 in the packing-LP framework of Garg and Koenemann as described in the proof of Theorem 4, we immediately get the following corollary:

Corollary 2. *There is an FPTAS for $\mathrm{BCMCFP}_{\mathbb{R}}$ that runs in $\mathcal{O}\left(\frac{1}{\varepsilon^2} \cdot m^2 \log m \cdot (\log m \log\log m + n \log n)\right)$ time on acyclic graphs.* □

5. Conclusion

In this paper, we presented results on the complexity and approximability of the budget-constrained minimum cost flow problem. As the problem is known to be solvable both in weakly polynomial time by interior-point methods and in strongly-polynomial time as shown in [5], we developed a new combinatorial algorithm that runs in weakly polynomial time $\mathcal{O}(\log M \cdot \mathrm{MCF}(m, n, C, U))$. Moreover, we presented a weakly polynomial-time FPTAS that uses the ε-approximate pareto frontier and a strongly polynomial-time FPTAS based on both Garg and Koenemann's packing-LP framework and Megiddo's parametric search technique. Moreover, we could show that we can improve the running-time of the latter algorithm for the case of acyclic graphs.

References

[1] R. K. Ahuja, T. L. Magnanti, J. B. Orlin, Network Flows, Prentice Hall, 1993.

[2] L. R. Ford, D. R. Fulkerson, Constructing Maximal Dynamic Flows from Static Flows, Operations Research 6 (1958) 419–433.
[3] V. Chankong, Y. Y. Haimes, Multiobjective Decision Making: Theory and Methodology, Dover Books on Engineering Series, Dover Publications, Incorporated, 2008.
[4] P. A. M. Duque, S. Coene, P. G. K., Sörensen, F. Spieksma, The accessibility arc upgrading problem, European Journal of Operational Research 224 (3) (2013) 458–465.
[5] M. Holzhauser, S. O. Krumke, C. Thielen, Budget-Constrained Minimum Cost Flows, Journal of Combinatorial Optimization 31 (4) (2016) 1720–1745.
[6] C. H. Papadimitriou, M. Yannakakis, On the approximability of trade-offs and optimal access of web sources, in: Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on, IEEE, 86–92, 2000.
[7] N. Garg, J. Koenemann, Faster and simpler algorithms for multicommodity flow and other fractional packing problems, SIAM Journal on Computing 37 (2) (2007) 630–652.
[8] N. Megiddo, Combinatorial optimization with rational objective functions, Mathematics of Operations Research 4 (4) (1979) 414–424.
[9] N. Megiddo, Applying parallel computation algorithms in the design of serial algorithms, Journal of the ACM (JACM) 30 (4) (1983) 852–865.
[10] E. Cohen, N. Megiddo, Maximizing concave functions in fixed dimension, World Scientific, 1990.
[11] S. Toledo, Approximate parametric searching, Information processing letters 47 (1) (1993) 1–4.
[12] S. Toledo, Maximizing non-linear concave functions in fixed dimension, in: Foundations of Computer Science, 1992. Proceedings., 33rd Annual Symposium on, IEEE, 676–685, 1992.
[13] A. Schrijver, Theory of Linear and Integer Programming, John Wiley & Sons, Chichester, 1998.
[14] P. M. Vaidya, Speeding-up linear programming using fast matrix multiplication, in: Foundations of Computer Science, 1989., 30th Annual Symposium on, IEEE, 332–337, 1989.
[15] M. Ehrgott, Multicriteria optimization, Springer, 2nd edn., 2005.
[16] A. M. Geoffrion, Solving bicriterion mathematical programs, Operations Research 15 (1) (1967) 39–54.
[17] L. G. Valiant, Parallelism in comparison problems, SIAM Journal on Computing 4 (3) (1975) 348–355.
[18] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, R. E. Tarjan, Linear time bounds for median computations, in: Proceedings of the fourth annual ACM symposium on Theory of computing, ACM, 119–124, 1972.
[19] R. M. Karp, A characterization of the minimum cycle mean in a digraph, Discrete mathematics 23 (3) (1978) 309–311.
[20] R. Chandrasekaran, Minimal ratio spanning trees, Networks 7 (4) (1977) 335–342.